Quantum Hyper-V plugin

Project blueprint

Author: **Alessandro Pilotti <apilotti@cloudbase.it>**
Version: **1.0**
Date: **01/10/2012**

Hyper-V reintroduction in OpenStack with the Folsom release was primarily focused on Nova compute support, limiting, due to time constraints, the networking features required in a modern cloud infrastructure to basic options not suitable for multi tenant and advanced scenarios.

The main goal of this project is to bring support for modern network configurations to deployments of any size, in particular multi tenant scenarios by creating a Quantum plugin that implements Hyper-V networking features.

The plugin will be implemented entirely in Python and released as open source (Apache 2.0) as part of the OpenStack Quantum project.

Interaction between Quantum and the underlying OS will be based on WMI, Powershell and if necessary Win32 API calls in a consistent way with the work done in the Nova and Cinder projects. Microsoft System Center VMM is not required.

Most of the implementation is straightforward and doesn't require particular discussion at this stage, with the notable exception of network isolation and security group implementations, as detailed in the following paragraphs.

Network isolation

Hyper-V supports the following options for network isolation.

Feature	Description	Min. Hyper-V version	Tenant identification in third party equipment
Network virtualization	This feature provides layer 3 encapsulation of the layer 2 data via NVGRE tunneling. The NVGRE IETF protocol is currently in draft status.	2012	Yes, via NVGRE protocol id
IP address rewrite	Each virtual machine customer address (CA) is mapped to a unique provider address (PA)	2012	No
VLAN	Traditional layer 2 virtual LANs	2008	Yes, via VLAN id

Each option has various pros and cons, depending on the deployment scenario. All three options should be implemented, letting the user choose which one suits better a given configuration for maximum flexibility.

Network virtualization

Also called Windows Network Virtualization (WNV), can be configured with the following Powershell cmdlets and underlying WMI calls:

- Enable-NetAdapterBinding
- Set-VMNetworkAdapter –VirtualSubnetId <id>
- New-NetVirtualizationProviderAddres
- New-NetVirtualizationCustomerRoute
- New-NetVirtualizationLookupRecord

Pros

- The number of IP and MAC addresses that need to be learned by third party devices decreases dramatically compared to other options due to the encapsulation.
- No specific configuration needed by third party devices, except for the NVGRE support.

Cons

- All the traffic in a given subnet is encapsulated in a single tunnel, which could bring performance issues in case of heavy traffic. In general limits related to any tunneling solution apply, e.g. loss of MTU optimizations / packet fragmentation.
- Tenant identification for QoS / traffic shaping purposes in third party equipment in the network topology (e.g. switches) requires hardware supporting NVGRE.
- QoS / traffic shaping based on CA addresses is harder due to the encapsulation and requires proper support.
- Optimal performances require NICs supporting GRE offloading on Hyper-V
- Almost non existing interoperability options at the moment in case of subnets spawning different hypervisors (e.g. Hyper-V, Xen, KVM)

IP address rewrite

Pros

- No need for an encapsulation protocol, each packet is simply rewritten with a different address (PA).
- QoS / traffic shaping can be performed, although there's no data in the IP address that can help in tying a packet to a tenant / VM beside the source address.

Cons

- A potentially large number of IP and MAC addresses need to be learned by third party hardware (e.g. switches).

VLAN

Pros

- VLANs are the traditional option in network isolation configuration, a tested, widely supported and well known solution.
- Excellent interoperability

Cons

- There's a maximum of 4096 possible values, which is a very tight limitation in large clouds.
- When a VLAN is added / modified, configurations need to be applied to a potentially large number of switches in the infrastructure, an error prone and complicated task, especially in case of heterogeneous hardware.
- A potentially large number of IP and MAC addresses need to be learned by third party hardware (e.g. switches).

Access control

Quantum defines the concept of “Security groups” to specify an access control list (ACL) based configuration for networking resources, mainly IP addresses and TCP / UDP ports.
From a feature perspective, it can be mainly considered as an extension of traditional firewalling and filtering solutions (e.g. iptables), brought outside of the VM and independent from the guest OS.

Hyper-V 2012 introduces a partial support for port ACLs, supporting only IP addresses and not TCP / UDP ports (*Add-VMNetworkAdapterAcl*, etc). Support for complete Quantum security groups on Hyper-V can be obtained by implementing an extensible switch driver based on the Windows Filter Platform (WFP) APIs.

Supported OS

- Windows Server 2012
- Hyper-V Server 2012